Information Flow Investigations: Extended Abstract*

Michael Carl Tschantz
mct@berkeley.edu
Univ. of California, Berkeley[†]

Anupam Datta
danupam@cmu.edu
Carnegie Mellon University

Jeannette M. Wing
wing@microsoft.com
Microsoft Research[†]

June 20, 2013

Nontraditional Information Flow Problems. Concerns about privacy have led to much interest in determining how third-party associates of first-party websites use the information they collect about the visitors to the first-party website. Some researchers have attempted to determine what these third-parties do with the information they collect [3, 6]. These researchers propose and use various analyses to determine what information is tracked and how it is used. They primarily design their analyses by intuition and do not formally present or study their analyses. Thus, questions remain: (1) Are the analyses used sound and/or complete? (2) Are they related to more formal prior work?

Furthermore, much work has been done on the detection of illicit flows of copyrighted files, such as work on *watermarking* [5, 4] and *traitor tracing* [1]. Similarly, companies handling sensitive data have adopted a variety of methods to discourage the misuse of such data by their employees. In particular, they employ watermarking-like counterintelligence operations to detect such leaks and determine the identity of the employee leaking the information [7].

In essence, each of the approaches used to solve these problems is an information flow analysis (IFA). In particular, the analyst would like to determine whether a system (a person or computer) is enabling a concerning flow of information. For example, an public advocacy group (an analyst) might want to determine whether Google (a system) uses a person's health-related web searches (a sensitive information source) to select advertisements (a low-level information sink).

However, despite the great deal of research on IFA, we know of no attempt to relate or inform the research on any of these problems with the models or techniques of IFA, even in an informal manner. We believe this disconnect exists since these problems differ from traditional IFA in an important respect. Whereas traditionally the analyst has access to the program running the analyzed system, in these problems, the analyst has no access to the program running the system, little control over its inputs, and a limited view of its behavior. Thus, the analyst does not have the information presupposed by traditional IFAs. To understand these problems as instances of IFA requires a fresh perspective on IFA.

Approach. Our goal is to systematize the information flow problems and analyses common to these areas of research. To do so, we identify and formalize the limited abilities of the analyst in these problems as the setting of *information flow investigations*, a form of analysis between the extremes of whitebox program analysis and blackbox monitoring (Fig. 1). We show that the ability of the analyst to control some inputs during an investigation enables powerful *sting* analyses that *setup* the system in question to discover its use of information without a whitebox model of the system. These analyses resemble the inductive reasoning used in experimental sciences. Our investigation framework provides a fresh perspective on our diverse set of motivating applications and allows us to elucidate and challenge approaches in these areas and in IFA.

*This research was supported by the U.S. Army Research Office grants DAAD19-02-1-0389 and W911NF-09-1-0273 to CyLab, by the National Science Foundation (NSF) grants CCF0424422 and CNS1064688, and by the U.S. Department of Health and Human Services grant HHS 90TR0003/01. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

[†] The authors conducted most of this work while at Carnegie Mellon University.

Contributions. In more detail, we formalize investigations in terms of a version of noninterference, the primary formal definition of traditional IFA [2]. We argue that every investigation is both unsound and incomplete for detecting (non)interference. Despite this negative result, we find our formalism useful for both explaining problems and solutions. In particular, we formalize problems from each motivating application and includes the first formal characterization of tracking web trackers. Presenting all three applications in terms of noninterference shows the relationship among these areas.

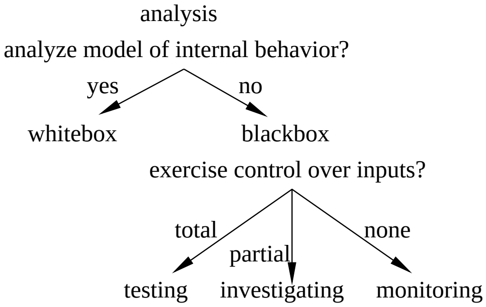

Figure 1: Taxonomy of analyses

Despite our unsoundness and incompleteness results, we identify a class of investigations we call *stings* that can produce strong guarantees under reasonable assumptions. They leverage the ability to control some inputs to the system to *setup* the system in such a way that its outputs reveals information use. The analyst employing these setups resembles a scientist manipulating factors during an experiment.

For each analysis used in the prior work on our motivating applications, we formalize it as a sting, which shows their similarities and differences. We derive the assumptions implicit in the informal analyses used in practice by studying them in our formalism as stings. These assumptions qualify the soundness and completeness of the conclusions drawn by works using stings. In particular, we produce practical suggestions for conducting future analyses of web trackers by applying our framework to prior works in the area.

Systematization of investigations is becoming increasingly important as technology trends (e.g., Cloud and Web services) result in analysts having limited access to and control over systems whose properties they are expected to study. Our work provides a useful starting point towards such a systematization by providing a common model and a shared vocabulary of concepts that ties together seemingly disparate areas of security and privacy by placing them in the context of information flow investigations.

Additional information about this ongoing research effort may be found at the following website:

`http://www.cs.cmu.edu/~mtschant/info-experiments/`

References

[1] Chor, B., Fiat, A., and Naor, M. Tracing traitors. In *Proceedings of the 14th Annual International Cryptology Conference on Advances in Cryptology* (1994), Springer-Verlag, pp. 257–270.

[2] Goguen, J. A., and Meseguer, J. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy* (1982), pp. 11–20.

[3] Guha, S., Cheng, B., and Francis, P. Challenges in measuring online advertising systems. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement* (2010), pp. 81–87.

[4] Swanson, M., Kobayashi, M., and Tewfik, A. Multimedia data-embedding and watermarking technologies. *Proceedings of the IEEE 86*, 6 (1998), 1064–1087.

[5] Wagner, N. R. Fingerprinting. In *Proceedings of the 1983 IEEE Symposium on Security and Privacy* (1983), p. 18.

[6] Wills, C. E., and Tatar, C. Understanding what they do with what they know. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society* (2012), pp. 13–18.

[7] Wright, P. *Spycatcher: The Candid Autobiography of a Senior Intelligence Officer.* Viking Adult, 1987.